A Bayesian Approach to Sparse plus Low rank Network Identification

Mattia Zorzi and Alessandro Chiuso

arXiv:1503.07340v2 [math.OC] 26 Sep 2015

***Abstract*— We consider the problem of modeling multivariate stochastic processes with parsimonious dynamical models which can be represented with a sparse dynamic network with few latent nodes. This structure translates into a sparse plus low rank model. In this paper, we propose a Bayesian approach to identify such models.**

I. INTRODUCTION

This paper deals with the network identification problem of a multivariate stochastic process of high dimension. Important applications can be found in many fields, for instance in econometrics, [1], [5], [8], [21], social network analysis [13], system biology, [20] and so on.

Let y be a stochastic process whose components (or variables) y_k are manifest, i.e. can be directly measured. We define as network of this manifest process y a directed graph wherein nodes denote the variables and edges encode conditional Granger causality relations among these variables, [12]. More specifically, there is an edge from node i (i.e. variable y_i) to node j (i.e. variable y_j) if the past of y_i is needed to predict y_j conditioned on the past of all the other y_k, $k \neq i$. Sparse graphs (i.e. with few edges) represent a concise and interpretable way to describe the relations among the variables of the manifest process, [7]. However, this modeling assumption does not exploit nor encode the fact that the manifest process may often be thought as driven by a low dimensional latent process, i.e. a stochastic process whose components cannot be directly measured.

In this paper we formulate a new network identification problem which takes into account the presence of this low dimensional latent process and the variables of the manifest process Granger causes each other mostly through the latent variables. The corresponding network has a two layer structure: one layer denotes the manifest variables and the other one the latent variables. The presence of few latent variables should drastically reduce the edges in the manifest layer, therefore increasing the degree of conciseness and robustness of the model. Finally, it turns out that a model described by this network has a sparse plus low rank (S+L) structure. More precisely, the low rank part depends on the number of latent variables, whereas the sparse one on the number of conditional Granger causality relations among the manifest variables.

This work has been partially supported by the FIRB project "Learning meets time" (RBFR12M3AC) funded by MIUR.

M. Zorzi is with the Dipartimento di Ingegneria dell'Informazione, Università degli studi di Padova, via Gradenigo 6/B, 35131 Padova, Italy zorzimat@dei.unipd.it

A. Chiuso is with the Dipartimento di Ingegneria dell'Informazione, Università degli studi di Padova, via Gradenigo 6/B, 35131 Padova, Italy chiuso@dei.unipd.it

When it comes to developing and identification algorithm, which maps measured data into an estimated dynamical model, it is fair to say that the prediction error method (PEM) is a consolidate paradigm in system identification [15], [19]. In the traditional setting, candidate models are described in fixed parametric model structures, e.g. ARMAX, whose complexity is determined using cross validation or information based criteria. Regularization has been recently introduced in the PEM framework, see [16], [17], [6], [18], as an alternative approach to control complexity of the estimated models. This latter class of methods start with a large enough (in principle infinite dimensional) model class; the inverse (ill-posed) problem of determining a specific model from a finite set of measured data can be made into a well posed problem using a penalty term, whose duty is to select models with specific features. In the Bayesian view, this is equivalent to the introduction of an *a priori* probability (i.e. prior) on the model to estimate. For instance, the prior should account the fact that the model is Bounded Input Bounded Output (BIBO) stable, [17].

The identification algorithm we propose belongs to this latter class of methods, where the predictor impulse response is estimated in the framework of Gaussian regression. The predictor impulse response is modeled as a zero mean Gaussian random vector. Its covariance matrix, referred to as kernel matrix, encodes the *a priori information*. In this case, the *a priori* information is that the predictor impulse responses are BIBO stable and that the model has a S+L structure. As we will see, such *a priori* information can be encoded in the kernel matrix using the maximum entropy principle. Moreover, this kernel matrix is characterized by the decay rate of the predictor impulse responses, by the number of conditional Granger causality relations among the manifest variables, and by the number of latent variables. These features are tuned by the so called hyperparameters vector. It is estimated by minimizing the negative log-likelihood of the measured data. Beside the fact that this problem is nonconvex, the joint estimation of the hyperparameters tuning the sparse and low rank part is not trivial, because these two parts may be nonidentifiable from the measured data. We propose an algorithm to estimate these hyperparameters which imposes and "hyper-regularizer" on the low rank hyperparameter to handle partially this non-uniqueness. Once the kernel matrix is fixed, an unique estimate of the S+L model is guaranteed through regularization.

We warn the reader that the present paper only reports some preliminary result regarding the Bayesian estimation of S+L models. In particular, all the proofs and most of the technical assumptions needed therein are omitted and will be

published afterwards.

The outline of the paper follows. In Section II, we introduce the S+L model for multivariate stochastic processes. In Section III, we apply the Gaussian approach to S+L system identification. In Section IV, we derive the kernel matrix by the maximum entropy principle. In Section V, we present the algorithm to estimate the hyperparameters of the kernel matrix. In Section VI, we provide some numerical examples to show the effectiveness of our method. Finally, the conclusions are in Section VII.

Notation

Throughout the paper, we will use the following notation. $\mathcal{S}_m^+$ denotes the cone of the positive definite symmetric matrices, and $\overline{\mathcal{S}}_m^+$ its closure. Given $v \in \mathbb{R}^m$ and $G \in \mathbb{R}^{n\times m}$, v_i denotes the i-th entry of v and $[G]_{ij}$ denotes the entry of G in position (i,j). $\|G\|_Q$ denotes the weighted Frobenius norm of G with weight matrix $Q \in \mathcal{S}_m^+$. Given a transfer matrix $L(z)$ of dimension $m \times m$, with some abuse of terminology, we say that $L(z)$ has rank equal to n, with $n \leq m$, if it admits the decomposition $L(z) = FH(z)$ where $F \in \mathbb{R}^{m\times n}$ and $H(z)$ is a $n \times m$ transfer matrix. Given a stochastic process $y = \{y(t)\}_{t\in\mathbb{Z}}$, with some abuse of notation, $y(t)$ will both denote a random vector and its sample value. Finally,

$$\mathrm{y}^-(t) := \begin{bmatrix} y(t-1)^\top & y(t-2)^\top & \dots \end{bmatrix}^\top \tag{1}$$

denotes the past data vector of y at time t. In similar way, $\mathrm{y}_i^-(t)$ denotes the past data vector of y_i.

II. S+L Models

Consider two zero mean stationary Gaussian stochastic processes y and x of dimension m and n, respectively. Let y be manifest, i.e. it can be measured, and x latent, i.e. it cannot be measured. We assume that $y(t) \in \mathbb{R}^m$ and $x(t) \in \mathbb{R}^n$ are described by the model

$$\begin{aligned} y(t) &= Fx(t) + S(z)y(t) + v(t) \\ x(t) &= H(z)y(t) + w(t) \end{aligned} \tag{2}$$

where $H(z) = \sum_{k=1}^\infty H_k z^{-k}$ is a BIBO stable transfer matrix of dimension $n \times m$, $F \in \mathbb{R}^{m\times n}$, $S(z) = \sum_{k=1}^\infty S_k z^{-k}$ is a BIBO stable transfer matrix of dimension $m\times m$, v and w are, respectively, m and n dimensional white Gaussian noise (WGN) with zero mean, and covariance matrix Σ_v and Σ_w. Moreover, v and w are independent.

It is possible to describe the structure of model (2) using a directed graph (i.e. network) as in Figure 1, [14]. Each node of this network corresponds to components of y or x; edges encode conditional Granger causality relations. More precisely, there is a direct link from node y_j (x_j) to node y_i (x_i) if and only if $\mathrm{y}_j^-(t)$ ($x_j(t)$) is needed to predict $y_i(t)$ ($x_i(t)$), conditionally on all the past information available at time t, i.e. $\mathrm{y}_k^-(t)$ with $k \neq j$ and $x(t)$ ($\mathrm{y}^-(t)$ and $x_k(t)$ with $k \neq j$). In this case, we shall say that y_j (x_j) conditionally Granger causes y_i (x_i), [12]. In Figure 1 we provide an example with $m = 6$ and $n = 1$. In particular, x_1 conditionally Granger causes all the components in y and conversely, and y_1 conditionally Granger causes y_5.

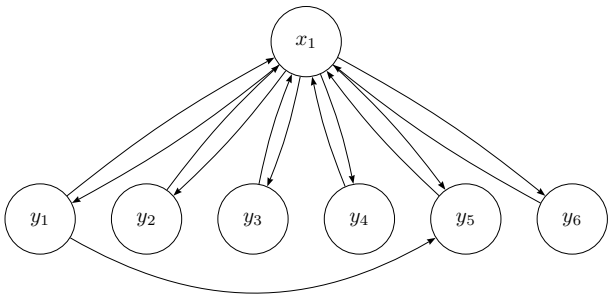

Fig. 1. Example of a sparse plus low rank model with $m = 6$ and $n = 1$.

Our main modeling assumption in (2) is that manifest variables Granger cause each other mostly through few latent variables. Therefore, we have $n \ll m$, i.e. the number of latent variables in x is small as compared to the number of manifest variables in y; similarly $S(z)$ is sparse, i.e. many of its entries are null transfer functions, so that few manifest variables conditionally Granger causes each other. From (2), we obtain the sparse plus low rank (S+L) model for y:

$$y(t) = S(z)y(t) + L(z)y(t) + e(t) \tag{3}$$

where $S(z)$ is a sparse transfer matrix by assumption, $L(z) := FH(z)$ is a low rank transfer matrix because F and $H^\top(z)$ are tall matrices, and $e(t) := v(t) + w(t)$ is WGN with covariance matrix $\Sigma = \Sigma_v + F\Sigma_w F^\top$.

Let $\hat{y}(t|t-1)$ be the minimum variance one-step ahead predictor of $y(t)$ based on the observations $\mathrm{y}^-(t)$, and $\hat{x}(t|t-1)$ be the minimum variance estimator of $x(t)$ based on $\mathrm{y}^-(t)$. From (2), we have

$$\begin{aligned} \hat{y}(t|t-1) &= F\hat{x}(t|t-1) + S(z)y(t) \\ \hat{x}(t|t-1) &= H(z)y(t). \end{aligned} \tag{4}$$

that is, the predictable part of y is a function of few estimated latent variables and of the pasts of few manifest variables of y. Eq. (4) can be compactly written as

$$\hat{y}(t|t-1) = S(z)y(t) + L(z)y(t). \tag{5}$$

It is worth noting that in the case that $L(z) = 0$, i.e. there is no need of latent variables to characterize the predictor of y, we obtain the sparse model presented in [7]. In the case $S(z) = 0$, i.e. the predictor of y is completely characterized by the estimators of the latent variables, we obtain a quasi-static factor model where the noise process is white, see for instance [10]. We would like to stress that the decomposition of a transfer matrix into sparse plus low rank may not be unique. As noticed in [5], this degeneracy may occur when $L(z)$ is sparse or $S(z)$ has few null entries. Although this nonidentifiability issue is important, the aim of this paper is to find one S+L decomposition (see Section III) which is not necessarily unique.

III. Gaussian Regression Approach to System Identification

Consider model (2) and assume that the measured data $y(1)\dots y(N)$ are extracted from a realization of y. The latent process x cannot be measured nor its dimension n is known. In this Section, we address the problem of estimating $S(z)$ and $L(z)$ from the given data. We draw inspiration from the Gaussian regression approach proposed in [16]. According this method, y is generated by model

$$y(t) = G(z)y(t) + e(t) \tag{6}$$

where $G(z)$ is a BIBO stable $m \times m$ transfer matrix and e is WGN with covariance matrix Σ. Note that, if we set $G(z) = S(z)+L(z)$ then (6) is equivalent to (3). On the other hand, with (6) we loose the S+L structure we are interest in.

Since $G(z)$ is BIBO stable, and thus the impulse response coefficients decay to zero as a function of the lag index, it is possible to use the approximation

$$G(z) = \sum_{k=1}^{T} G_k z^{-k} \tag{7}$$

where T is sufficiently large. The parameters G_k, $k = 1,.,T$ of the truncated transfer matrix are stacked in the vector $\theta \in \mathbb{R}^{m^2T}$ which is defined as follows

$$\begin{aligned}\theta \;=\; & \left[\; (g^{[11]})^\top \quad \dots \quad (g^{[1m]})^\top \;\middle|\; \dots \right. \\ & \left. \dots \;\middle|\; (g^{[m1]})^\top \quad \dots \quad (g^{[mm]})^\top \;\right]^\top\end{aligned} \tag{8}$$

where

$$g^{[ij]} = \left[\; [G_1]_{ij} \quad [G_2]_{ij} \quad \dots \quad [G_T]_{ij} \;\right]^\top \tag{9}$$

denotes the impulse response coefficients of the transfer function in position (i, j) of the truncated approximation of $G(z)$. Then, we stack the measured data in the vector y as follows

$$\begin{aligned}\mathrm{y} \;=\; & \left[\; y_1(T+1)^\top \quad \dots \quad y_1(N)^\top \;\middle|\; \dots \right. \\ & \left. \dots \;\middle|\; y_m(T+1)^\top \quad \dots \quad y_m(N)^\top \;\right]^\top. \end{aligned} \tag{10}$$

In similar way, we define

$$\begin{aligned}\mathrm{e} \;=\; & \left[\; e_1(T+1)^\top \quad \dots \quad e_1(N)^\top \;\middle|\; \dots \right. \\ & \left. \dots \;\middle|\; e_m(T+1)^\top \quad \dots \quad e_m(N)^\top \;\right]^\top. \end{aligned} \tag{11}$$

From (6) the vector of the measured data can be expressed in the linear regression form

$$\mathrm{y} = \Phi\theta + \mathrm{e} \tag{12}$$

where $\Phi \in \mathbb{R}^{mN\times m^2T}$ is the regression matrix. Note that, $\Phi\theta$ is the one-step ahead predictor of y.

According to the Gaussian regression framework, θ is modeled as a zero mean Gaussian random vector with covariance matrix, or kernel matrix, denoted by $K \in \mathcal{S}^+_{m^2T}$.

Let $\hat\theta$ be the posterior mean of θ given y. In [16] it has been proved that, under some technical assumption, $\hat\theta$ can also be written as solution to the Tikhonov regularization problem

$$\hat\theta = \arg\min_\theta \|\mathrm{y} - \Phi\theta\|^2_{\Sigma^{-1}\otimes I_N} + \|\theta\|^2_{K^{-1}}. \tag{13}$$

Moreover, it is not difficult to see that

$$\hat\theta = K\Phi^\top(\Phi K\Phi^\top + \Sigma\otimes I_N)^{-1}\mathrm{y}. \tag{14}$$

Remark 3.1: Using the theory of reproducing kernel Hilbert spaces, [3], it is possible to show that the above results still hold for $T \to \infty$, [16].

Remark 3.2: Although we assumed $K \in \mathcal{S}^+_{m^2T}$, the Bayes estimator (14) also holds for K singular. In that case, Problem (13) is well defined provided that θ belongs to the range of K.

The optimal solution $\hat\theta$ highly depends on the choice of K. A typical assumption is that the transfer functions in $G(z)$ are independent, that is $g^{[ij]}$ are independent vectors. We shall also assume that $g^{[ij]}$, with $i, j = 1\dots m$, are identically distributed, so that $K = I_{m^2} \otimes \tilde K$ where $\tilde K \in \mathcal{S}^+_T$ is the covariance matrix of $g^{[ij]}$, with $i, j = 1\dots m$. In this paper, $\tilde K$ is chosen as a filtered version of the tuned/correlated (TC) kernel, see [16] and [6] for more details. It is important to note this kernel is able to capture high frequency oscillations, which are typical of the predictor impulse responses for low pass processes, and enforces BIBO stability as $T \to \infty$ on the posterior mean of the predictor impulse responses.

A. Gaussian Regression Approach to S+L Identification

The idea is to model $S(z)$ and $L(z)$ through Gaussian process, similarly to what has been done above for $G(z)$. In particular, since $S(z)$ and $L(z)$ are BIBO stable, we can consider their truncated approximations with T is sufficiently large. Then, it is not difficult to see that

$$\mathrm{y} = \Phi(\theta_l + \theta_s) + \mathrm{e} \tag{15}$$

where y and e have been defined in (10) and (11). $\theta_s, \theta_l \in \mathbb{R}^{m^2T}$ contains the parameters of the truncated approximations of $S(z)$ and $L(z)$, respectively. Then, we model θ_s and θ_l as zero mean random vectors with covariance matrix K_S and K_L, respectively. Moreover, we shall assume that θ_s and θ_l are Gaussian and independent. As we will see in Section IV, these assumptions are suggested by the maximum entropy principle.

Proposition 3.1: Let $\hat\theta_s$ and $\hat\theta_l$ be, respectively, the posterior mean of θ_s and θ_l given y. Then, under some technical assumption, $\hat\theta_s$ and $\hat\theta_l$ are solution to the Tikhonov regularization problem

$$\arg\min_{\theta_s,\theta_l} \|\mathrm{y}-\Phi(\theta_s+\theta_l)\|^2_{\Sigma^{-1}\otimes I_N} + \|\theta_s\|^2_{K_S^{-1}} + \|\theta_l\|^2_{K_L^{-1}}. \tag{16}$$

Moreover, we have

$$\hat\theta_s = K_S\Phi^\top c, \;\; \hat\theta_l = K_L\Phi^\top c \tag{17}$$

where

$$c = (\Phi(K_S + K_L)\Phi^\top + \Sigma\otimes I_N)^{-1}\mathrm{y}. \tag{18}$$

In what follows, $\hat{\theta}_s$ and $\hat{\theta}_l$ will be referred to as posterior mean of $S(z)$ and $L(z)$, respectively. In the next Section, we shall show how K_S and K_L can be chosen so as to enforce BIBO stability on both the posterior mean of $S(z)$ and $L(z)$, sparsity on the posterior mean of $S(z)$ as well as low rank of the posterior mean of $L(z)$.

IV. Maximum Entropy Kernel Matrix

In this Section we characterize the prior probability density of θ_s and θ_l by using the maximum entropy principle. Such principle states that among all the prior probability densities satisfying certain desired constraints, the optimal one should maximize the differential entropy.

Our starting assumptions are that θ_s and θ_l are absolutely continuous zero mean random vectors. Let $\mathbf{p}(\theta_s, \theta_l)$ denote the joint probability density of θ_s and θ_l. Let $\mathbb{E}$ denote the integration over $\mathbb{R}^{2m^2T}$ with respect to the probability measure $\mathbf{p}$. Moreover, $\mathcal{P}$ denotes the space of probability densities which are Lebesgue integrable. The differential entropy of $\mathbf{p} \in \mathcal{P}$ is, [9],

$$\mathbf{H}(\mathbf{p}) = -\mathbb{E}[\log(\mathbf{p}(\theta_s, \theta_l))]. \tag{19}$$

Next, we characterize the constraint on θ_s enforcing BIBO stability and sparsity on the posterior mean of $S(z)$. The transfer function in position (i, j) of $S(z)$ is the null transfer function if and only if $s^{[ij]}$ is the null vector. We consider the constraint

$$\mathbb{E}[\|s^{[ij]}\|^2_{\tilde{K}^{-1}}] \leq p_{ij}, \tag{20}$$

where $p_{ij} \geq 0$. If $p_{ij} = 0$, then $s^{[ij]}$ is zero in mean square and so also is its posterior mean. Moreover, simple algebraic manipulations show that the weighted second moment bound in (20) implies a bound on the variance of k-th element of $s^{[ij]}$ which decays as the k-th element in the main diagonal of $\tilde{K}$. Therefore, condition (20) enforces BIBO stability and sparsity on the posterior mean of $S(z)$.

Regarding the low rank constraint on θ_l, let $A_l \in \mathbb{R}^{m \times mT}$ be the random matrix such that

$$A_l = \begin{bmatrix} L_1 & L_2 & \dots & L_T \end{bmatrix}. \tag{21}$$

Consider the constraint

$$\mathbb{E}[A_l(\tilde{K}^{-1} \otimes I_m)A_l^\top] \leq Q. \tag{22}$$

If $Q \in \overline{\mathcal{S}}_m^+$ has $m - n$ singular values equal to zero, then the posterior mean of $A_l A_l^\top$ has rank less than or equal to n. Therefore, the latter admits the decomposition

$$A_l = \begin{bmatrix} FH_1 & FH_2 & \dots & FH_T \end{bmatrix}, \tag{23}$$

where $F \in \mathbb{R}^{m \times n}$ and $H_k \in \mathbb{R}^{n \times m}$, $k = 1 \dots T$, as in Section II. Equivalently, the posterior mean of $L(z)$ admits the decomposition $L(z) = FH(z)$. Similarly to the sparse part, the weight matrix $\tilde{K}^{-1} \otimes I_m$ enforces BIBO stability on the posterior mean of $L(z)$.

Consider the following maximum entropy problem

$$\begin{aligned} \max_{\mathbf{p} \in \mathcal{P}} \quad & \mathbf{H}(\mathbf{p}) \\ \text{s.t.} \quad & \mathbb{E}[\|s^{[ij]}\|^2_{\tilde{K}^{-1}}] \leq p_{ij} \;\; i, j = 1 \dots m \\ & \mathbb{E}[A_l(\tilde{K}^{-1} \otimes I_m)A_l^\top] \leq Q \end{aligned} \tag{24}$$

where $p_{ij} > 0$ $i, j = 1 \dots m$, and $Q \in \mathcal{S}_m^+$.

Theorem 4.1: The optimal solution to (24) is such that θ_s and θ_l are independent, Gaussian with zero mean and covariance matrix

$$K_S = \Gamma \otimes \tilde{K}, \;\; K_L = \Lambda \otimes I_m \otimes \tilde{K} \tag{25}$$

where

$$\Gamma = \mathrm{diag}(\gamma_{11} \dots \gamma_{m^2}) \in \mathcal{S}_{m^2}^+ \tag{26}$$

and $\Lambda \in \mathcal{S}_m^+$.

The matrices Γ and Λ are the hyperparameters of the kernel matrices K_S and K_L, respectively. Clearly, we are interested in the limiting case where $p_{ij} = 0$ for some (i, j) and Q low-rank. Let $\mathbf{P} = \{(i, j) \text{ s.t } p_{ij} = 0\}$ and $\mathbf{Q} = \{v \in \mathbb{R}^m \text{ s.t } Qv = 0\}$. Then, it can be shown that the maximum entropy solution can be extended by continuity to this limiting case where $\gamma_{(i-1)m+j} = 0$ if and only if $(i, j) \in \mathbf{P}$ and $\Lambda v = 0$ if and only if $v \in \mathbf{Q}$. Thus, Γ tunes sparsity on the posterior mean of $S(z)$ and Λ tunes the rank on the posterior mean of $L(z)$. Finally, is worth noting that the hyperparameters tuning the decay rate of the posterior mean of the predictor impulse responses are encoded in $\tilde{K}$, [17].

V. Estimation of the Hyperparameters

In order to compute $\hat{\theta}_s$ and $\hat{\theta}_s$ we need to estimate the hyperparameters in $\tilde{K}$ and the matrices Γ and Λ. The hyperparameters describing $\tilde{K}$ are estimated in a preliminary step by minimizing the negative log-likelihood of model (6), see [16]. Γ and Λ are obtained minimizing the negative log-marginal likelihood ℓ of y. Under some technical assumption, we have, [16],

$$\ell(\mathrm{y}, \Gamma, \Lambda) = \frac{1}{2} \log \det V + \frac{1}{2} \mathrm{y}^\top V^{-1} \mathrm{y} + \text{const.term} \tag{27}$$

where

$$V = \Phi(K_S + K_L)\Phi^\top + \Sigma \otimes I_N. \tag{28}$$

Since (27) is nonconvex in V, only local minima can be computed. Beside that, the joint minimization of Γ and Λ is not trivial because the sparse and low rank part may be nonidentifiable from the measured data. For this reason, we constrain the structure of Λ as follows:

$$\Lambda = \alpha(I - UU^\top) + U\mathrm{diag}(\beta_1 \dots \beta_r)U^\top \tag{29}$$

where $U \in \mathbb{R}^{m \times r}$ and its columns are the first r singular vectors of an estimate $\hat{A}_l \hat{A}_l^\top$ of $A_l A_l^\top$. In this way, the constraints in Λ are decoupled along the "most reliable" r singular vectors of $\hat{A}_l \hat{A}_l^\top$ and their orthogonal complement. This is equivalent to fix r latent variables (from the estimate $\hat{A}_l \hat{A}_l^\top$). Regarding the hyperparameter Γ, in [2] it has been shown that the minimization of (27) leads sparsity in the main diagonal of Γ. Therefore, we minimize (27) with respect to $\tilde{\xi} = \{\gamma_1 \dots \gamma_{m^2}, \alpha, \beta_1 \dots \beta_r\}$ while r and U are fixed. The complete procedure to estimate r, U and $\tilde{\xi}$ is described in Algorithm 1. $r^{(k)}$, $U^{(k)}$, $\hat{A}_l^{(k)}$ and $\tilde{\xi}^{(k)}$ denote, respectively, r, U, $\hat{A}_l$ and $\tilde{\xi}$ at the k-th iteration. Finally, to minimize efficiently (27) with respect to $\tilde{\xi}$ we used the scaled gradient projection algorithm developed in [4].

Algorithm 1 Computation of r, U and $\tilde{\xi}$

$k = 0$
$r^{(0)} \leftarrow 0$
$U_{OPT}^{(0)} \leftarrow$ empty matrix
$\tilde{\xi}^{(0)} \leftarrow \underset{\tilde{\xi}}{\operatorname{argmin}}\, \ell(\mathrm{y}^+, \tilde{\xi}, U_{OPT}^{(0)})$
$\tilde{\xi}_{OPT}^{(0)} \leftarrow \tilde{\xi}^{(0)}$
repeat
  $k \leftarrow k + 1$
  $r^{(k)} \leftarrow r^{(k-1)} + 1$
  if $r^{(k)} = 1$ **then**
    $\hat{A}_l^{(k)} \leftarrow [\,\hat{G}_1 \;\; \hat{G}_2 \;\; \dots]$ where $\hat{G}_1, \hat{G}_2, \dots$ are the coefficients of $\hat{G}(z)$ estimated from (6) with $K = I_{p^2} \otimes \tilde{K}$
  else
    $\hat{A}_l^{(k)} \leftarrow [\,\hat{L}_1^{(k)} \;\; \hat{L}_2^{(k)} \;\; \dots]$ where $\hat{L}_1^{(k)}, \hat{L}_2^{(k)}, \dots$ are the coefficients of $\hat{L}^{(k)}(z)$ estimated from (16) with K_S and K_L having hyperparameters given by $U_{OPT}^{(r^{(k)}-1)}$ and $\tilde{\xi}_{OPT}^{(r^{(k)}-1)}$
  end if
  $U^{(k)} \leftarrow$ first $r^{(k)}$ singular vectors of $\hat{A}_l^{(k)}\ \hat{A}_l^{(k)\,\top}$
  $\tilde{\xi}^{(k)} \leftarrow \underset{\tilde{\xi}}{\operatorname{argmin}}\, \ell(\mathrm{y}^+, \tilde{\xi}, U^{(k)})$
  repeat
    $\tilde{\xi}_{OPT}^{(r^{(k)})} \leftarrow \tilde{\xi}^{(k)}$
    $U_{OPT}^{(r^{(k)})} \leftarrow U^{(k)}$
    $k \leftarrow k + 1$
    $r^{(k)} \leftarrow r^{(k-1)}$
    $\hat{A}_l^{(k)} \leftarrow [\,\hat{L}_1^{(k)} \;\; \hat{L}_2^{(k)} \;\; \dots]$ where $\hat{L}_1^{(k)}, \hat{L}_2^{(k)}, \dots$ are the coefficients of $\hat{L}^{(k)}(z)$ estimated from (16) with K_S and K_L having hyperparameters given by $U_{OPT}^{(r^{(k)}-1)}$ and $\tilde{\xi}_{OPT}^{(r^{(k)}-1)}$
    $U^{(k)} \leftarrow$ first $r^{(k)}$ singular vectors of $\hat{A}_l^{(k)}\ \hat{A}_l^{(k)\,\top}$
    $\tilde{\xi}^{(k)} \leftarrow \underset{\tilde{\xi}}{\operatorname{argmin}}\, \ell(\mathrm{y}^+, \tilde{\xi}, U^{(k)})$
  until $\ell(\mathrm{y}^+, \tilde{\xi}^{(k)}, U^{(k)}) < \ell(\mathrm{y}^+, \tilde{\xi}^{(k-1)}, U^{(k-1)})$
until $\ell(\mathrm{y}^+, \tilde{\xi}_{OPT}^{(r^{(k)})}, U_{OPT}^{(r^{(k)})}) < \ell(\mathrm{y}^+, \tilde{\xi}_{OPT}^{(r^{(k)}-1)}, U_{OPT}^{(r^{(k)}-1)})$
$r \leftarrow r^{(k)} - 1$
$U \leftarrow U_{OPT}^{(r^{(k)}-1)}$
$\tilde{\xi} \leftarrow \tilde{\xi}_{OPT}^{(r^{(k)}-1)}$

VI. Simulation results

We consider three Monte Carlo studies of 50 runs where at any run a model with $m = 6$ manifest variables is randomly generated. For each run in the Monte Carlo experiments an identification data set and a test set, both of size 500, are generated. The noise covariance matrix Σ is always estimated via a preliminary step using a low-bias ARX-model, see [11].

In the first experiment, the models have McMillan degree equal to 20, and are perturbed versions of (3) with $l = 1$ latent variable and four non null transfer functions in $S(z)$.

The second experiment is identical to the first one, with the exception that the latent variables are $l = 2$.

In the third experiment, the models have McMillan degree equal to 30, but without a special structure.

We compare the following one-step ahead predictors:

- TRUE: this is the one computed from the true model
- PEM: this is the one computed from the PEM approach, as implemented in pem.m function of the MATLAB System Identification Toolbox
- TC: this is the one computed with the approach described at the beginning of Section III
- SL: this is our method described in Section III-A.

The following performance indexes are considered:

- average relative complexity of the S+L network of the SL model (in percentage)

$$AC = \frac{100}{50} \sum_{t=1}^{50} \frac{\#SL_k}{m^2 T} \tag{30}$$

  where $\#SL_k$ is the number of parameters of the estimated S+L model at the k-th run, whereas m^2T is the number of parameters of a nonstructured model
- one-step ahead coefficient of determination (in percentage)

$$\mathrm{COD} = 100 \left(1 - \frac{\frac{1}{500} \sum_{t=1}^{500} \|y^{\mathrm{test}}(t) - \hat{y}^{\mathrm{test}}(t|t-1)\|^2}{\frac{1}{500} \sum_{t=1}^{500} \|y^{\mathrm{test}}(t) - \bar{y}^{\mathrm{test}}\|^2} \right)$$

  where $\bar{y}^{\mathrm{test}}$ denotes the sample mean of the test set data $y(1)^{\mathrm{test}} \dots y^{\mathrm{test}}(500)$ and $\hat{y}^{\mathrm{test}}(t|t-1)$ is the one-step ahead prediction computed using the estimated model
- average impulse response fit (in percentage)

$$\mathrm{AIRF} = 100 \left(1 - \frac{\sum_{k=1}^{50} \|G_k - \hat{G}_k\|^2}{\sum_{k=1}^{50} \|G_k - \bar{G}\|^2} \right) \tag{31}$$

  with $\bar{G} = \frac{1}{50} \sum_{k=1}^{50} G_k$.

TABLE I
Average relative complexity of the S+L Bayesian Network

Epx. #	#1	#2	#3
AC	55.89	63.72	81.56

Table I shows the percentage of the average relative complexity of the S+L network. In particular, in the first two experiments our method is able to detect that the underlying model is close to have a simple S+L network. Figure 2 shows the COD in the first two experiments. One can see that SL provides a slightly better performance than TC. On the other hand, SL provides better estimators for the predictor coefficients than the TC, Figure 3. Finally, Figure 4 shows the COD in the third experiment. The median of SL is slightly worse than the one of TC. On the other hand, the bottom whisker of SL is better than the one of TC. Indeed, SL simplified the S+L network, see Table I, increasing the robustness of the estimated predictor impulse response coefficients.

VII. Conclusions

In this paper, we proposed a Gaussian regression approach to identify multivariate stochastic processes having sparse network with few latent nodes. Simulations show that our approach is able to identify a S+L network which does not compromise the prediction performance.

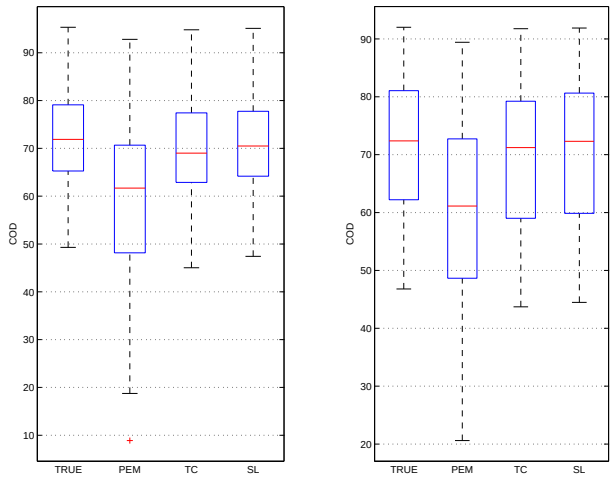

Fig. 2. One step ahead coefficient of determination in the first experiment (left panel) and in the second experiment (right panel).

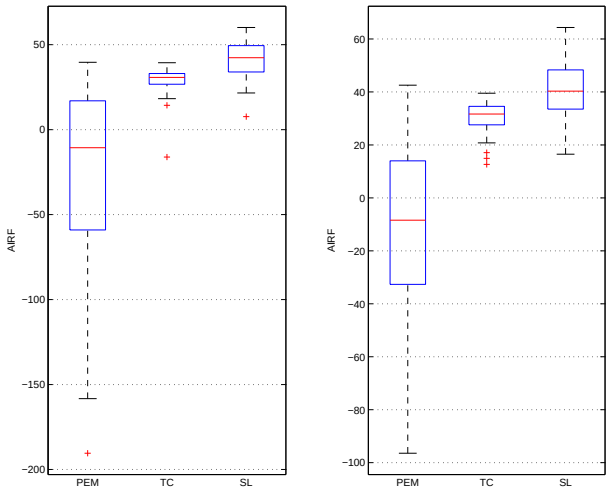

Fig. 3. Average impulse response fit in the first experiment (left panel) and in the second experiment (right panel).

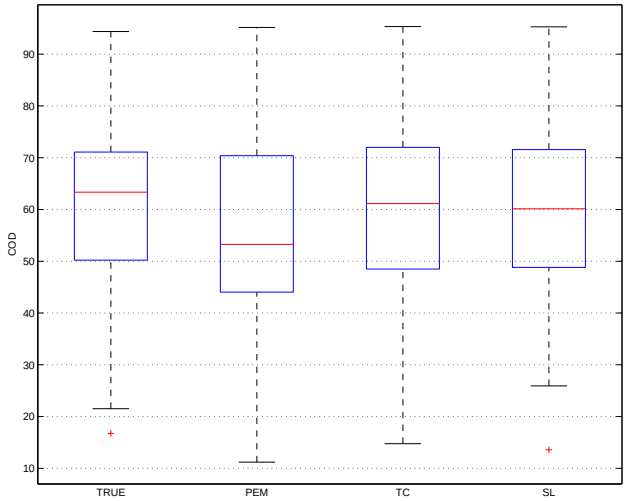

Fig. 4. One step ahead coefficient of determination in the third experiment.

REFERENCES

[1] A. Abdelwahab, O. Amor, and T. Abdelwahed. The analysis of the interdependence structure in international financial markets by graphical models. *Int. Res. J. Finance Econ.*, 15:291–306, 2008.

[2] A. Aravkin, J. Burke, A. Chiuso, and G. Pillonetto. Convex vs non-convex estimators for regression and sparse estimation: the mean squared error properties of ard and glasso. *Journal of Machine Learning Research*, 15:217–252, 2014.

[3] N. Aronszajn. Theory of reproducing kernels. *Trans. Amer. Math. Soc.*, 68:337–404, 1950.

[4] S. Bonettini, A. Chiuso, and M. Prato. A scaled gradient projection method for Bayesian learning in dynamical systems. *SIAM Journal on Scientific Computing*, 37:1297–1318, 2015.

[5] V. Chandrasekaran, P. Parrilo, and A. Willsky. Latent variable graphical model selection via convex optimization. *Annals of Statistics (with discussion)*, 40(4):1935–2013, Apr. 2010.

[6] T. Chen, H. Ohlsson, and L. Ljung. On the estimation of transfer functions, regularizations and gaussian processes-revisited. *Automatica*, 48(8):1525–1535, 2012.

[7] A. Chiuso and G. Pillonetto. A Bayesian approach to sparse dynamic network identification. *Automatica*, 48(8):1553–1565, 2012.

[8] M. Choi, V. Chandrasekaran, and A. Willsky. Gaussian multiresolution models: Exploiting sparse markov and covariance structure. *IEEE Transactions on Signal Processing*, 58(3):1012–1024, March 2010.

[9] T. Cover and J. Thomas. *Information Theory*. Wiley, New York, 1991.

[10] M. Deistler and C. Zinner. Modelling high-dimensional time series by generalized linear dynamic factor models: An introductory survey. *Communications in Information & Systems*, 7(2):153–166, 2007.

[11] G. Goodwin, M. Gevers, and B. Ninness. Quantifying the error in estimated transfer functions with application to model order selection. *IEEE Transactions on Automatic Control*, 37(7):913–928, Jul 1992.

[12] C. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, Aug. 1969.

[13] E Kolaczyk. *Statistical analysis of network data: methods and models*. Springer, 2009.

[14] S. Lauritzen. *Graphical Models*. Oxford University Press, Oxford, 1996.

[15] L. Ljung. *System Identification (2Nd Ed.): Theory for the User*. Prentice Hall, New Jersey, 1999.

[16] G. Pillonetto, A. Chiuso, and G. De Nicolao. Prediction error identification of linear systems: A nonparametric gaussian regression approach. *Automatica*, 47(2):291–305, 2011.

[17] G. Pillonetto and G. De Nicolao. A new kernel-based approach for linear system identification. *Automatica*, 46:81–93, 2010.

[18] G. Pillonetto, F. Dinuzzo, T. Chen, G. De Nicolao, and L. Ljung. Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, 50(3):657–682, 2014.

[19] T. Söderström and P. Stoica, editors. *System Identification*. Prentice Hall, New Jersey, 1988.

[20] A. Werhli, M. Grzegorczyk, and D. Husmeier. Comparative evaluation of reverse engineering gene regulatory networks with relevance networks, graphical gaussian models and Bayesian networks. *Bioinformatics*, 22(20):2523–2531, 2006.

[21] M. Zorzi and R. Sepulchre. AR identification of latent-variable graphical models. Conditionally accepted in IEEE Transactions on Automatic Control, 2015.